

Measuring the temperature and diversity of the U.S. regulatory ecosystem

Michael J Bommarito II^{a,b}, Daniel Martin Katz^{a,b}

^a*Illinois Tech - Chicago Kent College of Law*

^b*CodeX - The Stanford Center for Legal Informatics*

Abstract

Over the last 23 years, the U.S. Securities and Exchange Commission has required over 34,000 companies to file over 165,000 annual reports. These reports, the so-called “Form 10-Ks,” contain a characterization of a company’s financial performance and its risks, including the regulatory environment in which a company operates. In this paper, we analyze over 4.5 million references to U.S. Federal Acts and Agencies contained within these reports to build a mean-field measurement of temperature and diversity in this regulatory ecosystem, where companies are organisms inhabiting the regulatory environment. While individuals across the political, economic, and academic world frequently refer to trends in this regulatory ecosystem, far less attention has been paid to supporting such claims with large-scale, longitudinal data. In this paper, we document an increase in the regulatory energy per filing, i.e., a warming “temperature.” We also find that the diversity of the regulatory ecosystem has been increasing over the past two decades, as measured by the dimensionality of the regulatory space and distance between the “regulatory bitstrings” of companies. These findings support the claim that regulatory activity and complexity are increasing, and this measurement framework contributes an important step towards improving academic and policy discussions around legal complexity and regulation.

Keywords: complex systems, natural language, temperature, diversity, legal complexity, regulation

1. Introduction

Economies, like ecosystems, exhibit dynamic, complex behaviors resulting from the interaction of “organisms” inhabiting and altering their “environment.” In the case of economies, organisms can be seen as companies, and environments can be seen, at least in part, as regulations. Just as changes in the environment like rising tem-

perature can harm or help organisms, either broadly or for specific regions or organisms, so too can regulation harm or help companies. Yet unlike studies of biological ecosystems, studies of the economy have thus far lacked a longitudinal, empirical measure of fundamental environment factors like “temperature” or “species diversity.” In this paper, we attempt to bridge this gap, finding

support for the common claim that regulatory activity and complexity has increased over the last 20 years.

Each year, companies that meet the registration requirements under the Securities and Securities Exchange Acts of 1933 and 1934 must file an annual report with the U.S. Securities and Exchange Commission (“SEC”). This report, the so-called Form 10-K as defined in sections 13 and 15(d) of the Securities Exchange Act of 1934 (15 U.S.C. § 78o or 78o(d)), provides a broad overview of a company’s performance and its risks. Unlike other sources of information, the statements contained within these reports are certified and attested to by both a company’s officers and its independent auditors. If these statements are negligent or fraudulent, the SEC, Department of Justice, and shareholders are all able to press civil and, in limited cases, criminal charges against officers and auditors.

The law encourages and sometimes requires disclosure, but the incentives for firms are not entirely one-sided. Countering the trend towards limitless disclosure is the competition for investor capital. Through their annual reports, companies generally seek to present themselves as better investment opportunities than their competitors. Their officers thus face a balancing act as between describing a pessimistic future full of potential risks and an optimistic future without any risks. While certainly not a perfect description of reality, no other source of information, including surveys and press releases, provides as comprehensive of a statement of the environment in which companies operate.

While 10-Ks contain a significant amount of information regarding risks facing the company, we are particularly interested in the regulatory risk and uncertainty high-

lighted in these filings. Thus, in this paper, we analyze more than 20 years, 30,000 companies, and 160,000 10-K reports to identify more than 4.5 million references to U.S. Federal regulatory Acts and Agencies. Using these references, we generate a reproducible, quantitative, and longitudinal measurement of the temperature and diversity of the U.S. regulatory “ecosystem.” We document a clear trend towards increasing total energy, temperature, and diversity, with double- to triple-digit growth in all measurements. We believe this framework and its ongoing application represent a principled approach to the quantification of both the global and various local regulatory ecosystems, and we hope that this research can drive better-grounded discussions of legal complexity and policy design in the modern world.

2. Data

10-K filings have been the focus of many academic studies in finance and accounting [1], [2], [3], [4], law [5], [6] and other adjacent fields [7], [8], [9]. Many of these studies have focused upon questions such as whether the reporting requirements are achieving their intended purpose and the extent to which markets react to the information disclosed in such findings. While these are certainly worthy questions in their own right, we believe these filings reveal important patterns that, *at scale* and *over time*, provide meaningful insight into a range of other scientific questions.

Companies that meet the requirements for 10-K reporting expend significant and increasing resources to prepare these documents, typically with the assistance of accounting firms and lawyers. Indeed, the Annual Audit Fee Survey [10] conducted by the

Financial Executives Research Foundation reveals a mean and median 2015 expense of \$1.8M and \$522,205, respectively, across over 6,000 filers. As required by law, the figures and statements contained within these reports are certified and attested to by both a company’s officers and its independent auditors. Therefore, unlike other sources of information, in general, these 10-K annual reports are more likely to convey comprehensive and realistic description of the environment in which companies operate.

Form 10-K filings generally contain at least four parts and fifteen schedules, which collectively offer a wealth of useful information about registered companies. These parts include a characterization of a company’s financial health, legal risks, and other systematic and idiosyncratic factors, such as the nature of the regulatory environment in which it operates. Some of these factors, such as tax credits, may be positive, but the majority of listed regulatory factors are present as risks. While there are a number of specific requirements under the law, firms and industries are provided with some latitude regarding how to satisfy reporting requirements. In addition, as explored in [1] and [5], there have been some important changes in the reporting rules over time. That said, SEC form templates and accounting firm standards result in more similarity than difference across firms and across time.

Through this exercise of risk factor disclosure, companies typically describe various sources of regulatory risk, including the laws and administrative agencies that are most relevant to their respective businesses. Consider, for example, the 2015 10-K filing of Trans Energy, Inc., an oil and gas exploration company. Among other statutory and agency references, their filing references

both the Migratory Bird Treaty Act of 1918 and Endangered Species Act of 1973:

The Endangered Species Act (“ESA”) was established to protect endangered and threatened species. Pursuant to that act, if a species is listed as threatened or endangered, restrictions may be imposed on activities adversely affecting that species’ habitat. Similar protections are offered to migratory birds under the **Migratory Bird Treaty Act**. The Company conducts operations on oil and natural gas leases that have species, such as raptors, that are listed and species, such as sage grouse, that could be listed as threatened or endangered under the ESA.

Across the broader set of required company disclosures, 10-K filings are filled with references such as these; some, such as citations to sections 13 and 15(d) of the Securities Exchange Act of 1934, are boilerplate and required by the SEC’s forms. Others, such as the 756 references to the Migratory Bird Treaty Act of 1918, or the 27 references to the Price-Anderson Nuclear Industries Indemnity Act of 1957 over the last 23 years, are not.

In the mid-1990s, the SEC introduced its Electronic Data Gathering, Analysis, and Retrieval (“EDGAR”) system. Since then, nearly all registered company 10-K reports have been uploaded and made available on EDGAR, resulting in more than 160,000 10-K reports accessible online. We retrieve these 10-K reports and build a multi-stage pipeline that identifies and normalizes references to Acts and Agencies. References

are first identified through standard natural language processing techniques; once a reference fragment is identified, it is then passed through a second stage of normalization. As one example, many filers reference the “Gramm Leach Bliley Financial Services Modernization Act of 1999;” however, they do not do so using its full name, as above. Instead, they frequently refer to it as “GLB,” “Graham Leach Bliley,” “Gramm Leach Bliley,” or the “Financial Services Modernization” Act. In order to handle this variation, we built a mapping for over 600 potential Act references, relying on a combination of the US Code, Wikipedia, and manual review. This mapping is then combined with fuzzy-string matching techniques to correct for spelling mistakes such as “Graham Leach Bliley.” The result is a high-precision and high-recall extraction of 401 unique Federal Acts and 133 Agencies across our 23-year dataset. In total, we identify more than 4.5 million Act and Agency references contained in 10-K reports over the past 23 years.

3. Methodology and Results

Temperature, at its theoretical basis, is a measure of energy per unit of area or volume. The more energy per unit of space, *c.p.*, the higher the temperature. While the actual behavior of matter may vary based on whether this energy is vibrational, rotational, or translational, temperature *per se* provides an incredibly valuable mean-field characterization of the behavior thermodynamic system.

In our context, each reference to an Act or Agency is equivalent to some energy expended by an accountant, attorney, or company to describe some dimension of regulatory exposure. While not every poten-

tial exposure a company faces is listed, and those that are listed do not have an “equal” energy that can be measured in J or W , these counts overall correspond to the total number and range of regulatory concepts addressed therein. This therefore allows us to leverage these references to provide a mean-field characterization of the overall regulatory environment inhabited by registered companies.

Like the global mean temperature of the Earth itself, there are certainly localized “regulatory micro-climates” whose temperature differs from the global climate. We appreciate that these micro-climates are of great importance across many professional and academic questions, but for purposes of this paper, we focus on aggregate, longitudinal patterns in this paper, relying on millions of references in hundreds of thousands of filings for tens of thousands of companies across more than two decades.

To operationalize these ideas, we first define a regulatory “space” and filing “profile.” The regulatory space consists of 401 dimensions - one for each discrete Act currently identified in our data. We can represent a filing as a vector in this regulatory space. In the simplest form, a 10-K profile $\vec{p}(a)$ has element $p_i(a)$ equal to the number of references to Act i for a company a ’s annual filing. This vector \vec{p} can then be normalized or projected. For example, we can project \vec{p} from the number of references to a “bitstring” vector $\vec{b}(a)$, whose element $b_i(a)$ is equal to 1 if a company a ’s filing mentions Act i at least once, else 0. Alternatively, we can aggregate or normalize these filings by viewing them as a time-indexed collection. Let $F(t)$ be the matrix whose rows correspond to the set of all \vec{p} or \vec{b} vectors filed within a given year t . Then the grand sum

of $F(t)$ is the total number of references, the j^{th} column sum of $F(t)$ is the number of references to the j^{th} Act, and the i^{th} row sum is the number of references per firm in year t . It is possible to normalize the number of references per Act to a rate of reference per filing by dividing the column sums of $F(t)$ by the number of rows $m(t)$, which we call $\bar{r}_j(t)$ for j^{th} Act in year t .

We then use these representations to measure the total energy, temperature, and diversity of the regulatory ecosystem as follows. First, we measure the total energy of the regulatory ecosystem using p -vectors as:

$$E(t) = \sum_i \sum_j F_{i,j}(t) \quad (1)$$

Figure 1 shows that the total energy, as measured by number of references to Acts per year, has increased substantially in the last 23 years. In 1996, there were just over 40,000 references to Acts in the nearly 5,000 filings that year; by 2006, these numbers had more than quadrupled to nearly 200,000 references in just over 9,000 filings; and, through three quarters of 2016, these numbers have again increased to an annualized rate of over 300,000 references.

Total energy alone can be misleading with respect to policy interpretations, however, as there are a number of reasons why energy may change without relation to regulatory exposure or “burden.” For example, i) the economy may grow or shrink in real or nominal terms, increasing or decreasing the number of companies or companies meeting registration requirement, ii) the SEC rules governing registration or filings may change, increasing or decreasing the number of companies or references, or iii) *c.p.*, the relative incentives to incorporate or take on shareholders may change, increasing or decreasing the number of companies registered.




Figure 1: Total number of Act references per filing over time.

These factors do not necessarily imply more or less regulation as experienced by individual companies, although they may be viewed as endogenous to some policy questions.

We can control for these factors by normalizing total energy to temperature, taking into account the number of filings per year as an analogy to area or volume. To do so, we calculate the average rate of references per filing, “temperature,” $T(t)$ as follows:

$$T(t) = \frac{E(t)}{m(t)} \quad (2)$$

Figure 2 shows that, over the last 23 years, $T(t)$ has been monotonically increasing. While the rate of reference, like the total energy in Figure 1 above, clearly shows the effect of the Sarbanes-Oxley changes in 2003, this trend remains unbroken both well before and well after. In 1996, the average number of references per filing was 8.4; by 2006, it had more than doubled to 20.9; and by 2016, the rate had increased again by more than 50% to 31.7 Act references per filing. Even if the amount of energy or cost does not scale linearly per filing with the

number of references, the monotonic, 237% increase in $T(t)$ clearly demonstrates an increasing regulatory temperature.




Figure 2: Average number of Act references per filing over time.

Table 1 below summarizes the data from Figures 1 and 2 above.

Year	$E(t)$	$T(t)$
1995	17,672	7.9
2000	75,851	10.6
2005	166,518	18.2
2010	227,210	25.0
2015	242,107	30.8

Table 1: Summary of energy and temperature measures over time.

Finally, we may ask - is temperature or energy changing in concert with diversity, or is the change in temperature concentrated along a single dimension of regulation? For example, changes in energy or temperature can represent more or less reliance on the same Act, e.g., the Securities Exchange Act of 1934; in this case, the number of unique Acts referenced is not changing, but the regulatory exposure per Act is. Alternatively, the total number of unique Acts referenced could be increasing or decreasing;

for example, in 2003, most registered companies added references to the recently enacted Sarbanes-Oxley Act, which had not previously been referenced. Changes such as these represent an increase or decrease in the number of dimensions or diversity of regulatory exposure, but not necessarily the intensity of each exposure.

Using our notation above, we evaluate the diversity question by calculating two measures. First, we calculate the number of unique Acts per filing through the sum of b vectors above. Then, we calculate the average number of unique Acts per filing, across all companies in a given year; this is $\frac{1}{m} \sum_j \sum_i b_i^j$, where b_i^j is the bit corresponding to whether the i^{th} Act was referenced in the j^{th} company filing and m is the number of companies per year.

Figure 3 shows that, over the last 23 years, the diversity of Acts referenced has increased jointly with temperature. Like Figures 1 and 2 above, the time series exhibits a jump following Sarbanes-Oxley; however, like Figure 2, the time series also exhibits a monotonic increase over two decades, growing from 3.1 unique Acts per filing in 1996 to 5.6 unique Acts per filing in 2006 to 7.9 Acts per filing in 2016. This increase suggests that the increase in regulatory ecosystem temperature has been, at least in part, related to an increase in the number of regulatory “micro-climates” present in the global regulatory ecosystem.

As an additional measure of diversity, we analyze each company’s yearly regulatory “bitstring.” As noted earlier, we calculate the 401-bit vector b for each company-year, where each bit corresponds to the presence of the 401 discrete Acts we identify. Although the regulatory space has 401 dimensions, the bitstring for a given filing is




Figure 3: Average number of unique Acts and Agencies per filing over time.

likely to be extremely sparse. For example, consider the 2012 10-K filed by the *Boeing Company*. Their filing features a bitstring with 12 non-zero elements, including Acts such as the Homeland Security Act, the Employee Retirement Income Security Act, the Patient Protection And Affordable Care Act and the American Taxpayer Relief. Alternatively, the 2014 10-K of *Facebook Inc.* features 10 unique elements, including The Bank Secrecy Act, the U.S. Foreign Corrupt Practices Act, the USA Patriot Act, and the Credit CARD Act.

After applying this formalization to all companies for all years, we calculate the average pairwise Hamming distance [11] between all company bitstrings in a given year. Hamming distance is commonly used to evaluate the diversity of genomic [12], [13], [14] and other related data [15], [16], [17]. It can be interpreted as proportional to the average number of regulatory dimensions not in common between companies. More explicitly, the Hamming distance between two companies a and b in year t is:

$$d_{a,b}(t) = \sum_i \vec{p}(a) \oplus \vec{p}(b) \quad (3)$$

where \oplus is the element-wise XOR operator. We can then write the average Hamming distance $\bar{d}(t)$ as the average over all combinations of a and b at t . Figure 4 visualizes the structure of the distance matrix D for all a, b as of 1994. The large block in the lower right corresponds primarily to special purpose vehicles like trusts or limited partnerships, and the overall structure corresponds to sectors and industries.




Figure 4: Hamming distance heatmap for 1994 Act bitstrings.

Figure 5, the average Hamming distance $\bar{d}(t)$ over time for Acts and Agencies, portrays mean-field distance between firms at scale, confirming an increasing diversity across the global regulatory ecosystem. Over time, companies are subject, on average, to increasingly different requirements. While not monotonically increasing like the rate of reference and number of unique references above, the average distance increases 18 of 23 years in the sample. In 1996, two firms were separated on average by fewer than four regulatory Act “bits” or “genes”;

by 2016, this number has increased to nearly 10.




Figure 5: Average pairwise Hamming distance between company filing regulatory bitstrings over time.

Table 2 summarizes the data from Figures 3 and 4 above.

Year	Average Unique Acts	$\bar{d}(t)$
1995	2.9	3.7
2000	3.6	4.5
2005	5.0	6.1
2010	6.2	6.8
2015	7.5	8.9

Table 2: Summary of diversity measures over time.

4. Conclusion and Future Work

In this paper, we have presented the first large-scale, longitudinal characterization of the energy, temperature, and diversity of the regulatory ecosystem. We have identified increasing regulatory exposure along an increasing number of dimensions, providing evidence in support of the claim that regulatory burden is increasing. Using a bitstring representation of firm regulatory exposure,

we have confirmed that the aggregate Federal regulatory ecosystem is becoming more diverse over time, providing evidence in support of the claim that regulatory complexity is increasing. These conclusions are based on more than 20 years, 30,000 companies, 160,000 10-K reports, and 4.5 million references contained in uniquely comprehensive and accurate 10-K reports.

In future work, we intend to expand upon these questions and connect to extant research agenda, including the categorization of regulatory “species” and “climates” through their time series signatures and linguistic markers, a mapping of the time-dependent, tri-partite network of firms, Acts, and Agencies, and the integration of this analysis with our existing work on the complexity of other US statutory, regulatory, and judicial systems [18], [19].

Our work contributes to both the broader literature on legal complexity [20], [21], [22] and efforts to document the physical properties of legal systems as complex adaptive systems [23], [24], [25], [26], [27], [28]. In addition, this paper is among a growing set of recent works applying tools of machine learning and natural language processing to better understand the behavior of various legal systems [29], [30], [31].

In sum, we believe that this framework for representation and measurement will contribute to ongoing academic and policy discussions around legal complexity and policy design. The continued development of both global and local regulatory indices can provide for a principled, empirical basis of evaluation, standing in stark contrast to some of the vague generalizations that frequently guide current discourse.

References

- [1] J. L. Campbell, H. Chen, D. S. Dhaliwal, H.-m. Lu, L. B. Steele, The information content of mandatory risk factor disclosures in corporate filings, *Review of Accounting Studies* 19 (2014) 396–455.
- [2] H. You, X.-j. Zhang, Financial reporting complexity and investor underreaction to 10-k information, *Review of Accounting Studies* 14 (2009) 559–586.
- [3] T. Kravet, V. Muslu, Textual risk disclosures and investors’ risk perceptions, *Review of Accounting Studies* 18 (2013) 1088–1122.
- [4] F. Li, The information content of forward-looking statements in corporate filings—a naïve bayesian machine learning approach, *Journal of Accounting Research* 48 (2010) 1049–1102.
- [5] K. K. Nelson, A. C. Pritchard, Carrot or stick? the shift from voluntary to mandatory disclosure of risk factors, *Journal of Empirical Legal Studies* 13 (2016) 266–297.
- [6] K. L. Doran, E. L. Quinn, Climate change risk disclosure: a sector by sector analysis of sec 10-k filings from 1995–2008, *North Carolina Journal of International Law and Commercial Regulation* 34 (2009) 101–147.
- [7] Y. Bao, A. Datta, Simultaneously discovering and quantifying risk types from textual risk disclosures, *Management Science* 60 (2014) 1371–1391.
- [8] K.-W. Huang, Z. Li, A multilabel text classification algorithm for labeling risk factors in sec form 10-k, *ACM Transactions on Management Information Systems (TMIS)* 2 (2011) 1–19.
- [9] M.-F. Tsai, C.-J. Wang, On the risk prediction and analysis of soft information in finance reports, *European Journal of Operational Research* 257 (2017) 243–250.
- [10] F. E. R. Foundation, 2016 audit fee report, 2016.
- [11] R. W. Hamming, Error detecting and error correcting codes, *Bell System technical journal* 29 (1950) 147–160.
- [12] D. C. Tully, C. B. Ogilvie, R. E. Batorsky, D. J. Bean, K. A. Power, M. Ghebremichael, H. E. Bedard, A. D. Gladden, A. M. Seese, M. A. Amero, et al., Differences in the selection bottleneck between modes of sexual transmission influence the genetic composition of the hiv-1 founder virus, *PLoS Pathog* 12 (2016) e1005619.
- [13] M. X. He, S. V. Petoukhov, P. E. Ricci, Genetic code, hamming distance and stochastic matrices, *Bulletin of mathematical biology* 66 (2004) 1405–1421.
- [14] C. Mattiussi, M. Waibel, D. Floreano, Measures of diversity for populations and distances between individuals with highly reorganizable genomes, *Evolutionary Computation* 12 (2004) 495–515.
- [15] E. Martin, E. Cao, Euclidean chemical spaces from molecular fingerprints: Hamming distance and hempel’s ravens, *Journal of computer-aided molecular design* 29 (2015) 387–395.
- [16] V. Alfred, Algorithms for finding patterns in strings, Elsevier, 2014.
- [17] R. W. Morrison, K. A. De Jong, Measurement of population diversity, in: *International Conference on Artificial Evolution (Evolution Artificielle)*, Springer, pp. 31–41.
- [18] M. J. Bommarito, D. M. Katz, A mathematical approach to the study of the united states code, *Physica A: Statistical Mechanics and its Applications* 389 (2010) 4195–4200.
- [19] D. M. Katz, M. J. Bommarito II, Measuring the complexity of the law: the united states code, *Artificial Intelligence and Law* 22 (2014) 337–374.
- [20] J. Ruhl, D. M. Katz, Measuring, monitoring, and managing legal complexity, *Iowa Law Review* 101 (2015) 191–244.
- [21] R. Boulet, P. Mazzega, D. Bourcier, Network analysis of the french environmental code, in: *Ai approaches to the complexity of legal systems. complex systems, the semantic web, ontologies, argumentation, and dialogue*, Springer, 2010, pp. 39–53.
- [22] D. Bourcier, P. Mazzega, Toward measures of complexity in legal systems, in: *Proceedings of the 11th international conference on Artificial intelligence and law*, ACM, pp. 211–215.
- [23] E. D. Lee, C. P. Broedersz, W. Bialek, Statistical mechanics of the us supreme court, *Journal of Statistical Physics* (2015) 1–27.
- [24] T. S. Clark, B. E. Lauderdale, The genealogy of law, *Political Analysis* 20 (2012) 329–350.
- [25] D. M. Katz, D. K. Stafford, Hustle and flow: A social network analysis of the american federal judiciary, *Ohio State Law Journal* 71 (2010)

- 457–507.
- [26] T. A. Smith, The web of the law, San Diego L. Rev. 44 (2007) 309–354.
 - [27] G. T. Jones, Dynamical jurisprudence: law as a complex system, Ga. St. UL Rev. 24 (2007) 873–883.
 - [28] D. G. Post, M. B. Eisen, How long is the coastline of law? thoughts on the fractal nature of legal systems, Journal of Legal Studies 29 (2000) 545–584.
 - [29] N. Aletras, D. Tsarapatsanis, D. Preotiu-Pietro, V. Lampos, Predicting judicial decisions of the european court of human rights: A natural language processing perspective, PeerJ Computer Science 2 (2016) e93.
 - [30] D. M. Katz, M. J. Bommarito, J. Blackman, A general approach for predicting the behavior of the supreme court of the united states, arXiv preprint arXiv:1612.03473 (2016).
 - [31] B. B. McShane, O. P. Watson, T. Baker, S. J. Griffith, Predicting securities fraud settlements and amounts: a hierarchical bayesian model of federal securities class action lawsuits, Journal of Empirical Legal Studies 9 (2012) 482–510.